

MOTION-BASED SYSTEM IDENTIFICATION AND FAULT DETECTION AND ISOLATION TECHNOLOGIES FOR THRUSTER CONTROLLED SPACECRAFT¹

Edward Wilson - NASA Ames/Intellization, Moffett Field, CA²
David W. Sutter - NASA Ames/Intellization
Dustin Berkovitz - MIT Dept. of Aeronautics/Astronautics, Cambridge, MA
Bradley J. Betts - NASA Ames/CSC
Edmund Kong - MIT Dept. of Aeronautics/Astronautics
Rommel del Mundo - NASA Ames/QSS Group, Inc.
Christopher R. Lages - NASA Ames/Intellization
Robert W. Mah - NASA Ames Research Center³
Richard Papsin - NASA Ames Research Center

ABSTRACT

By analyzing the motions of a thruster-controlled spacecraft, it is possible to provide on-line (1) thruster fault detection and isolation (FDI), and (2) vehicle mass- and thruster-property identification (ID). Technologies developed recently at NASA Ames have significantly improved the speed and accuracy of these ID and FDI capabilities, making them feasible for application to a broad class of spacecraft. Since these technologies use existing sensors, the improved system robustness and performance that comes with the thruster fault tolerance and system ID can be achieved through a software-only implementation. This contrasts with the added cost, mass, and hardware complexity commonly required by FDI. Originally developed in partnership with NASA - Johnson Space Center to provide thruster FDI capability for the X-38 during re-entry, these technologies are most recently being applied to the MIT SPHERES experimental spacecraft to fly on the International Space Station in 2004. The model-based FDI uses a maximum-likelihood calculation at its core, while the ID is based upon recursive least squares estimation. Flight test results from the SPHERES implementation, as flown aboard the NASA KC-135A 0-g simulator aircraft in November 2003 are presented.

INTRODUCTION

Thruster-controlled spacecraft (vehicles, probes, satellites, etc.) can benefit from system identification (ID) and fault detection and isolation (FDI) systems. The FDI system, control system, state-estimation system, image processing, and other systems on-board or on the ground can benefit from accurate estimates of mass and thruster properties. FDI systems often use additional sensors such as pressure sensors in the thruster nozzles; however when building and operating smaller, less expensive spacecraft, there is a call for intelligent fault tolerance without requiring increased hardware complexity, cost, and mass.

The algorithms presented here for FDI and mass- and thruster-property identification (ID) use existing navigation sensors such as gyros or accelerometers. Therefore, these added capabilities could be provided as a software-only addition to a new spacecraft or modification to an existing one. These algorithms can be implemented on-board using the spacecraft's existing processor(s), on-board using a stand-alone processor, or on the ground, processing sensor information communicated to the ground stations from the spacecraft.

¹ Portions of this research were performed with funding from NASA Headquarters, HQ AA, PWC 349-00, and by the NASA Intelligent Systems Program 302-10-10 (part of the CICT Program).

² Ed.Wilson@intellization.com

³ Robert.W.Mah@nasa.gov, group lead of the Smart Systems Research Lab in the Computational Sciences Division (code IC) at NASA Ames Research Center

The algorithms have been developed through application in MATLAB software simulation [MATLAB 2003] to four space vehicles as shown in Figure 1: the X-38 vehicle 201⁴ (space-based test vehicle for the CRV, Crew Return Vehicle) [X38 2003] and Mini-AERCam (Mini-Autonomous Extravehicular Robotic Camera) [AERCam 2003], both developed at NASA Johnson Space Center; the NASA Ames Research Center Smart Systems Research Lab air-bearing vehicle (S4), and the MIT SPHERES (Synchronized Position Hold, Engage, Reorient, Experimental Satellites) experimental spacecraft to fly on the International Space Station [SPHERES 2003]. They are presently being implemented in hardware on the S4 and SPHERES. The same software is used on all four vehicles (except for the real-time SPHERES code which is a C-implementation of the algorithms), indicating the generic nature of the technology and its implementation.

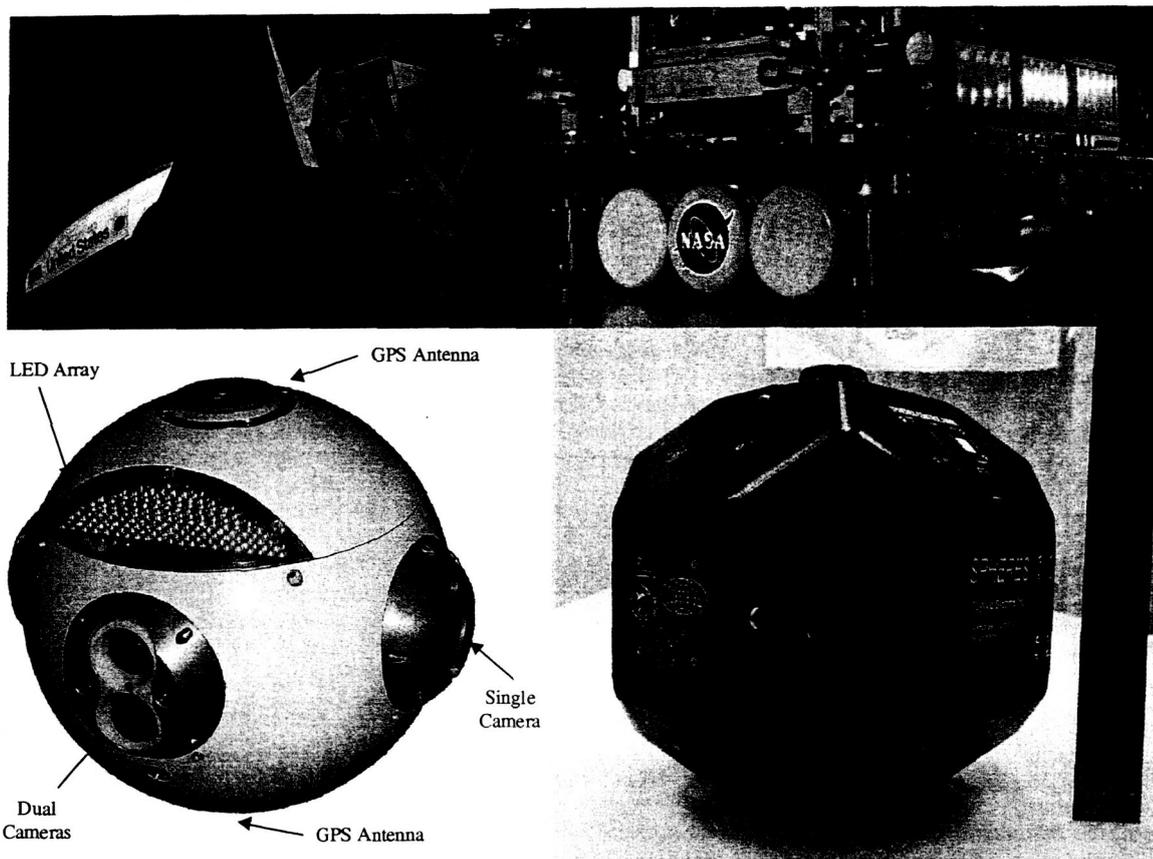


Figure 1: Four of the vehicles successfully tested in simulation. Clockwise from upper left: X-38 v. 201 with de-orbit propulsion stage; NASA Ames S4 air bearing vehicle; MIT SPHERES; JSC Mini AERCam. Hardware implementation and testing is currently underway on the S4 and SPHERES.

This paper summarizes the algorithms, and reports on tests of the mass-property ID algorithms on the MIT SPHERES as flown on NASA's zero-g KC-135A in November 2003.

RELATED RESEARCH – FAULT DETECTION AND ISOLATION

Several FDI approaches reported in the literature [Isermann 1984] perform well on a variety of applications. However, the on-off nature of the thrusters present in the class of applications addressed here limits the viability of many general-purpose methods. For example, if a thruster has failed off, it will appear to be working correctly at all times that it is not commanded to fire. This paper presents a general approach for this class of problems that has been validated through application to specific, realistic spacecraft applications.

Deyst and Deckert [Deyst 1976] at MIT/Draper Lab developed a maximum-likelihood based approach for detecting leaking thrusters for the Space Shuttle orbiter's RCS jets. The method for detecting soft failures was also

⁴ This since-canceled program was the original driver of the FDI basic research and development. The mass-property ID development was originally driven by a desire to provide accurate mass parameters for use by the FDI system.

extended to detect hard RCS jet failures. The maximum-likelihood method presented in that work is used and extended in this research.

Wilson and Rock [Wilson 1995a] [Wilson 1995b] at Stanford University developed an FDI method based on exponentially weighted recursive least squares estimation using accelerometer and angular-rate sensors. A neural network then provided adaptive control reconfiguration to multiple destabilizing hard and soft thruster failures. This was experimentally demonstrated on a 3-degree-of-freedom air-bearing vehicle.

Lee and Brown [Lee 1998] at JPL developed a leak monitoring system for the Cassini spacecraft that incorporated a model of the spacecraft.

RELATED RESEARCH – MASS-PROPERTY IDENTIFICATION

The basic difficulty in on-line mass property identification is that the mass parameters (center of mass, products and moments of inertia, mass) do not all appear linearly in the system's equations of motion. The solution approach presented here, and taken by others, is to make assumptions regarding the system uncertainties, motions, disturbances, and available computational power, and then manipulate these equations into forms that can be addressed using existing techniques such as RLS or second order filters.

Tanygin and Williams [Tanygin 1997] developed a least squares (LS) based algorithm to identify mass properties for a spinning vehicle during coasting maneuvers.

Bergmann, et al. [Bergmann 1987] [Bergmann 1990] [Richfield 1988] developed an ID approach using a Gaussian second-order filter [Gelb 1974], which resembles an extended Kalman filter, but has extra terms to address the second order effects. This is significantly more complex and computationally intensive (by about two orders of magnitude) than the approach presented here, and may not produce better results for most spacecraft. Due to the extra complexity, it may be more susceptible to noise and parameter variations than the presented methods. It assumes perfect knowledge of thruster properties.

Wilson and Rock [Wilson 1995a] [Wilson 1995b] developed an ID method based on exponentially weighted RLS using accelerometer and angular rate sensors. The acceleration created by each thruster (reflecting both mass and thruster properties) was identified. A neural network then provided adaptive control reconfiguration to multiple destabilizing hard and soft thruster failures. This approach (identifying thruster acceleration rather than separately identifying mass and thruster properties) is more direct (since thruster acceleration is the real value of interest from a control, estimation, or FDI standpoint), and probably better for vehicles with properties that are truly unknown (such as for the case where deflected thrusters are allowed, as on the vehicle tested in that research). However, for most vehicles, certain properties are well known, such as the thruster directions and locations in the structural frame, or the rate of fuel mass expulsion. The approach presented here can take advantage of that knowledge to get better estimates of the properties that are not well known.

PROBLEM DEFINITION

These technologies are designed for application to thruster controlled spacecraft equipped with motion sensors sufficient to provide an accurate estimate of angular acceleration. The FDI and ID performance will vary with the accuracy of this estimate. In the spacecraft studied here, gyros are used, but use of other sensors (video, star tracker, etc.) is possible as long as angular acceleration can be estimated and related to the corresponding thruster firing period. The failure mode types studied are listed below. The maximum likelihood FDI presented here requires a finite list of possible modes to choose from, although leak detection technology has also been developed to address soft failures.

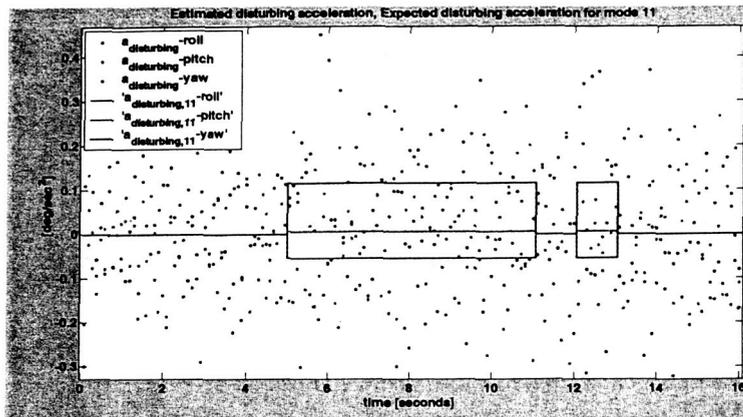
Both technologies are designed to run in a purely passive mode (with no controller intervention), but if control inputs can be guided, better and faster ID and FDI is possible.

THRUSTER FDI AND MASS-PROPERTY ID TECHNOLOGIES

The basic algorithms underlying these technologies have been published [Wilson 2002a] [Wilson 2002b]. Specifics on the major areas of development are as follows:

MOTION-BASED THRUSTER FDI

A maximum-likelihood-based approach to thruster FDI for spacecraft was developed. The system uses gyro signals (and accelerometers if desired) to detect and isolate hard, abrupt single- and multiple-jet on- and off-failures. Faults are detected within one second and identified within one to five seconds in most cases for the X-38. In extended testing for the X-38, failures were correctly identified in 99.98% of the test cases. A detailed description of the FDI algorithm is presented in [Wilson 2002a].



Although the FDI operates fundamentally by using the seemingly trivial comparison of expected accelerations and estimated accelerations, the major challenge is doing this rapidly, passively, and accurately in the presence of significant sensor noise, temporal thruster variability, imperfectly known vehicle mass and thruster properties, and multiple fault modes having similar acceleration signatures. This figure shows the estimated disturbing angular acceleration (blue, green, red dots at each control update for each axis), along with the expected disturbing acceleration for the true failure mode (solid lines, becoming non-zero when thruster 11 is commanded to fire). The

SNR is clearly very low in this case, and successfully isolating the failure mode is very challenging. This is an example of what necessitated the development of the full FDI system, in contrast to and as an extension of the simpler maximum-likelihood core developed by Deyst and Deckert [Deyst 1976].

The motion-based thruster FDI technology was developed to optimize speed and accuracy for the X-38 application. Since that was an exceptionally difficult application, once developed for the X-38, application to the other spacecraft examples has been very simple, requiring only setting updated parameters (thruster locations, etc.) and detection thresholds for the new spacecraft examples.

MOTION-BASED SYSTEM ID

Least-squares algorithms were developed to identify the vehicle center of mass, inertia matrix, inverse inertia matrix, and the force produced by each thruster. These algorithms have been implemented in both batch and recursive implementations. A novel approach to system ID was developed in which multiple recursive least squares (RLS) IDs run concurrently to ID parameters that appear as multiples of one another in the governing equations. The efficiency and robustness of this approach was judged to be sufficiently beneficial to offset the lack of theoretical optimality as might be offered by a higher order filter or nonlinear optimization. It was judged that the remaining unmodeled deviations from nominal (e.g., in thruster directions) would prevent the improved accuracy offered by those more complex approaches from actually being achieved. A detailed description of the ID algorithms is presented in [Wilson 2002b]. Some extensions since that publication are discussed briefly here.

SPHERES implementation

The Mass ID algorithms have been implemented in Embedded C++ for execution on the SPHERES TI TMS320C6701 floating point digital signal processor (DSP) on a Sundance SMT375 board. The code sends down telemetry for post-flight analysis, including raw gyro integers, filtered rates and angular accelerations, and ID results. The performance of the real-time code has been verified to match results of the original MATLAB code within precision expected with floating point processing.

Angular acceleration estimation

The on-board angular rate and acceleration estimator processes the 1 kHz gyro data corresponding to each control-update period (100 ms in this case). Since these results are also used for thruster FDI, it was decided that these filters should not assume any knowledge of the thrusters other than that the acceleration should be constant during the entire control-update period. So, for example, an abrupt, hard thruster failure could occur between one control segment and the next. Also driven by the need for on-line FDI, these filters must provide outputs in real-time, so non-causal smoothing filters were not considered. Although still under refinement, this estimator, whose example results are shown later, successfully provides accurate estimates in the presence of thruster latency and often-significant gyro ringing while requiring a relatively small amount of compute time and program memory space.

Inertia matrix identification

In addition to the inverse-inertia matrix identification, inertia matrix identification is now performed as well. Since the initial SPHERES zero-g experiments have shown a not-insignificant gyroscopic term ($\omega \times (I\omega)$), this term is now incorporated directly into the identification, rather than being handled as a calculated disturbance as had been suggested before.

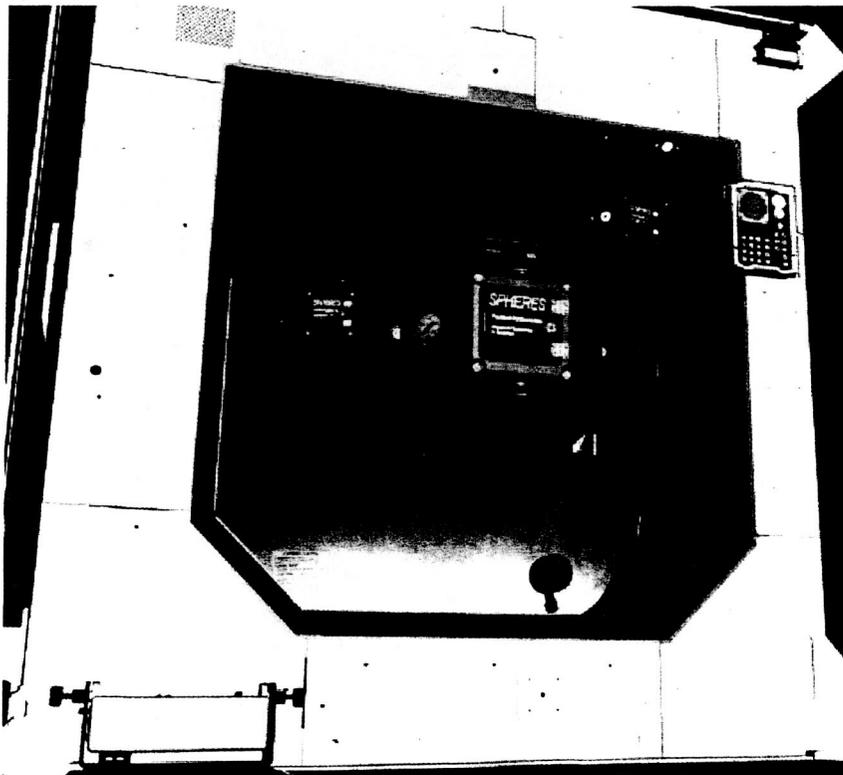
Thruster bias force identification

Another identification added is the average force produced by each thruster (one parameter for each thruster). Since thruster direction is relatively more easily characterized before flight, thruster directions are still treated as known values.

Thruster system characterization

Ground-based testing of the SPHERES spacecraft has progressed towards developing accurate characterization of several thruster properties, including: mass-flow rate, as a function of the number of jets firing; the multi-jet scale factor that represents the reduction in thrust for each thruster when multiple jets are fired simultaneously.

RESULTS AND DISCUSSION



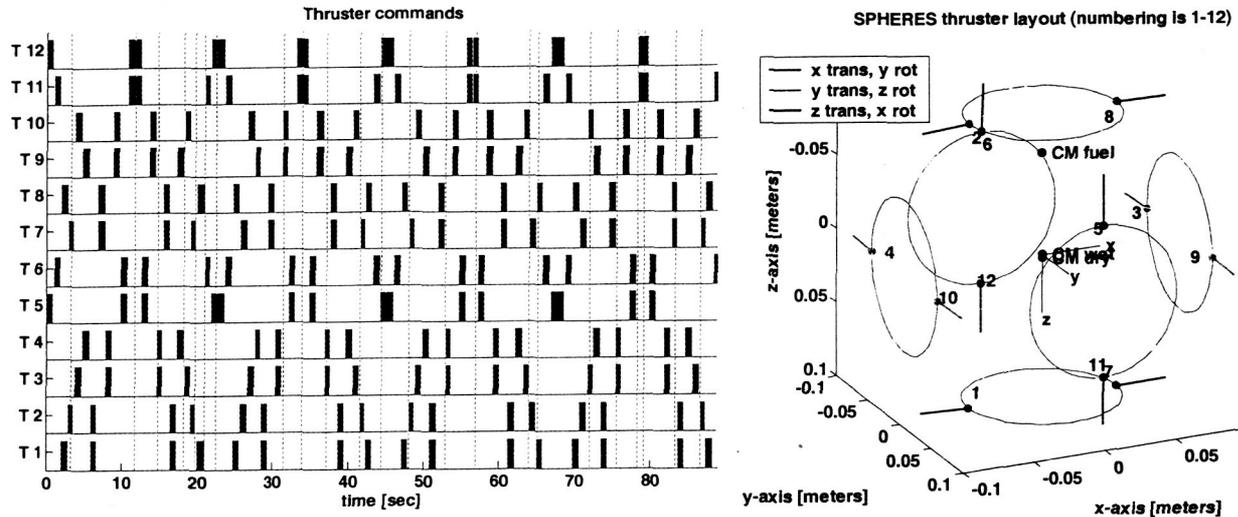
The thruster FDI and mass-property ID algorithms are being implemented on the SPHERES hardware for upcoming flight tests aboard the International Space Station (ISS). The MIT Space Systems Laboratory and Payload Systems, Inc. (PSI) are leading development of the SPHERES experimental spacecraft and coordinating its launch to the ISS, which is expected in 2004. The NASA Ames/Intellization team is working with MIT and PSI to integrate the algorithm implementations onto the spacecraft. A snapshot of a 3-D visualization of a 3-SPHERES experiment is shown here. In addition to the three SPHERES flying the ISS US Lab, ultrasonic beacons used by the global position and attitude determination system (PADS) can be seen in the lower left and upper right corners.

In a near-final test prior to launch, the mass ID algorithm was tested in a series of KC-135A ("Vomit Comet") flights from November 5-8, 2003, providing zero-g testing to complement the extensive 1-g air-bearing table testing performed at MIT.

The experimental results presented here were conducted on 8 November 2003 by Dustin Berkovitz and others from the MIT/PSI team, using SPHERES #5. Nine consecutive parabolas (nominally 20 seconds of zero-g, but practically, closer to about 10 seconds of free-floating time) were allocated for the Mass ID experiments. Data from other flights also proved useful, but the flight presented here contains the single longest set of data.

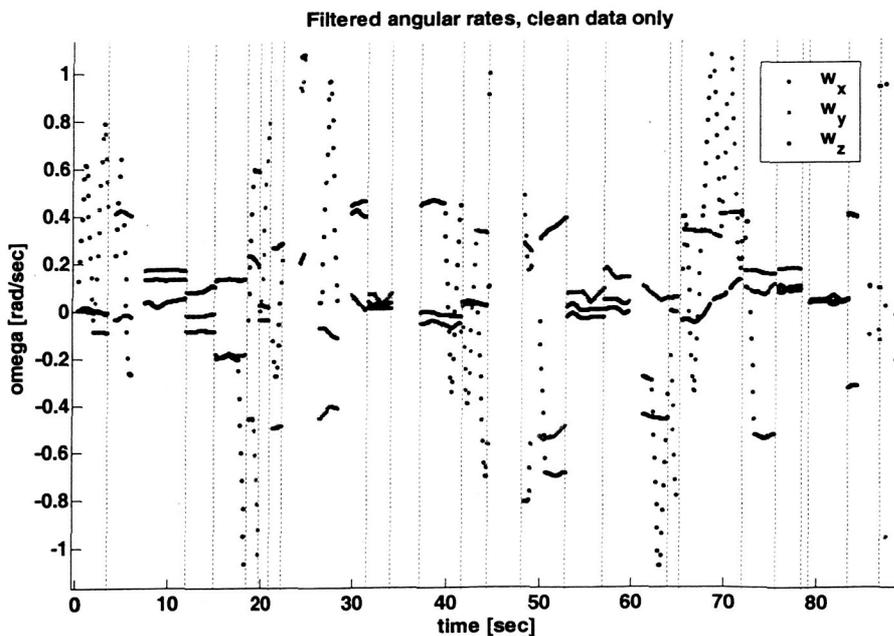
Through use of a "pause" feature in the control software, these 9 parabolas were all part of a single experiment consisting of about 90 seconds of zero-g flight time, with the system paused when handling was required and between parabolas. To test the response of the ID, a proof mass was attached using Velcro and removed between parabolas. To determine the effect of fuel slosh, a fuel tank pre-weighed at 28g was swapped in between parabolas as well. This mass approximately corresponds to full pressure (5.9 MPa = 860 psi) with no liquid CO₂, so no fuel slosh should have been present for those parabolas.

Although the mass ID algorithms ran on-board, due to disturbances when the SPHERES was handled (that occurred inadvertently when the system was un-paused), the analysis presented here is based on processing the telemetered rates and accelerations using the original MATLAB code, after cleaning the data to remove detected disturbances (handling to prevent contact with the walls of the aircraft) and periods of gyro saturation. The real-time software presently detects gyro saturation and does not update the ID during those times, but it does not presently try to detect handling.

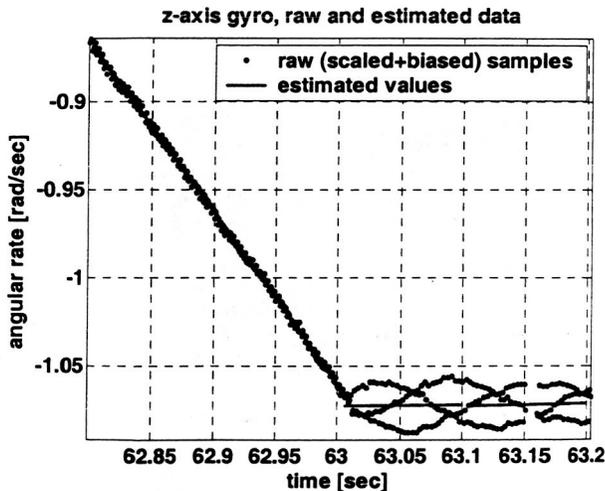
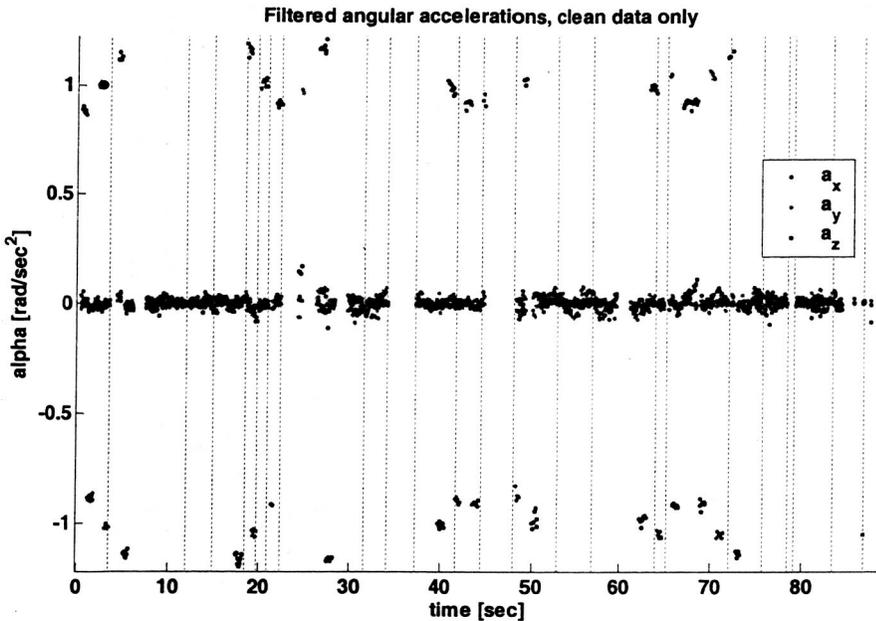


Since PADS was non-operational for the KC flights, the 12 SPHERES thrusters were commanded in repeating, 24-second long, open loop sequence as shown here. There were always either 0 or 2 thrusters firing to minimize the effect of uncertainty in thrust reduction with multiple thrusters firing. This sequence fires forward then backward on each of the 3 rotational axes, then the 3 translational axes, then all that backwards. Each pulse is 0.8 seconds long,

followed by 0.2 seconds of coasting. These lengths were chosen to maximize the firing time without causing excessive gyro ringing. Although the ID algorithm can accommodate arbitrary firing patterns, this simple sequence was chosen to facilitate manual checking. Since some of the skipped telemetry samples are not shown here, the first repetition of the sequence begins at 22.4 seconds rather than 24.0. In this and many of the plots that follow, times when the system was paused are indicated with a vertical green line.

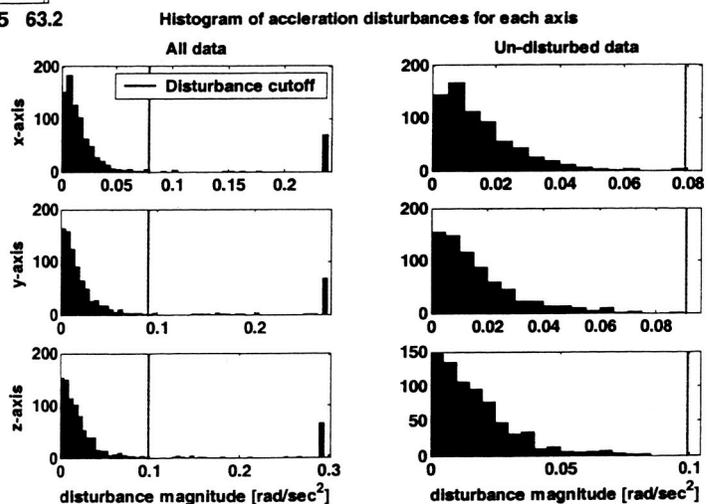


These figures show angular rates and accelerations for all 9 parabolas, after post-flight cleaning to remove



previous plot. The segments defined by the on-board-estimated rates and accelerations are overlaid on the raw gyro data to demonstrate the estimator's ability to filter out the gyro ringing while providing accurate and high bandwidth estimation. The acceleration estimates (slopes of the red lines in the plot) before and after thruster closing are: $[-0.9735 \ -0.9665 \ -0.9773 \ 0.0133 \ 0.0167]$. This raw data will be analyzed to further refine the on-board estimation algorithm.

In cleaning the data, the estimated acceleration at each sample was compared to that expected, based upon the best estimate of mass properties, and the thrusters fired. This



data during handling or gyro near-saturation (where the output is nonlinear). This data is the telemetered output of the on-board rate and acceleration estimator. The non-zero acceleration around 24 seconds is due to the gyroscopic term $(\omega \times (I\omega))$, which was larger than had been expected due to the relatively high rates on multiple axes and asymmetry of the SPHERES.

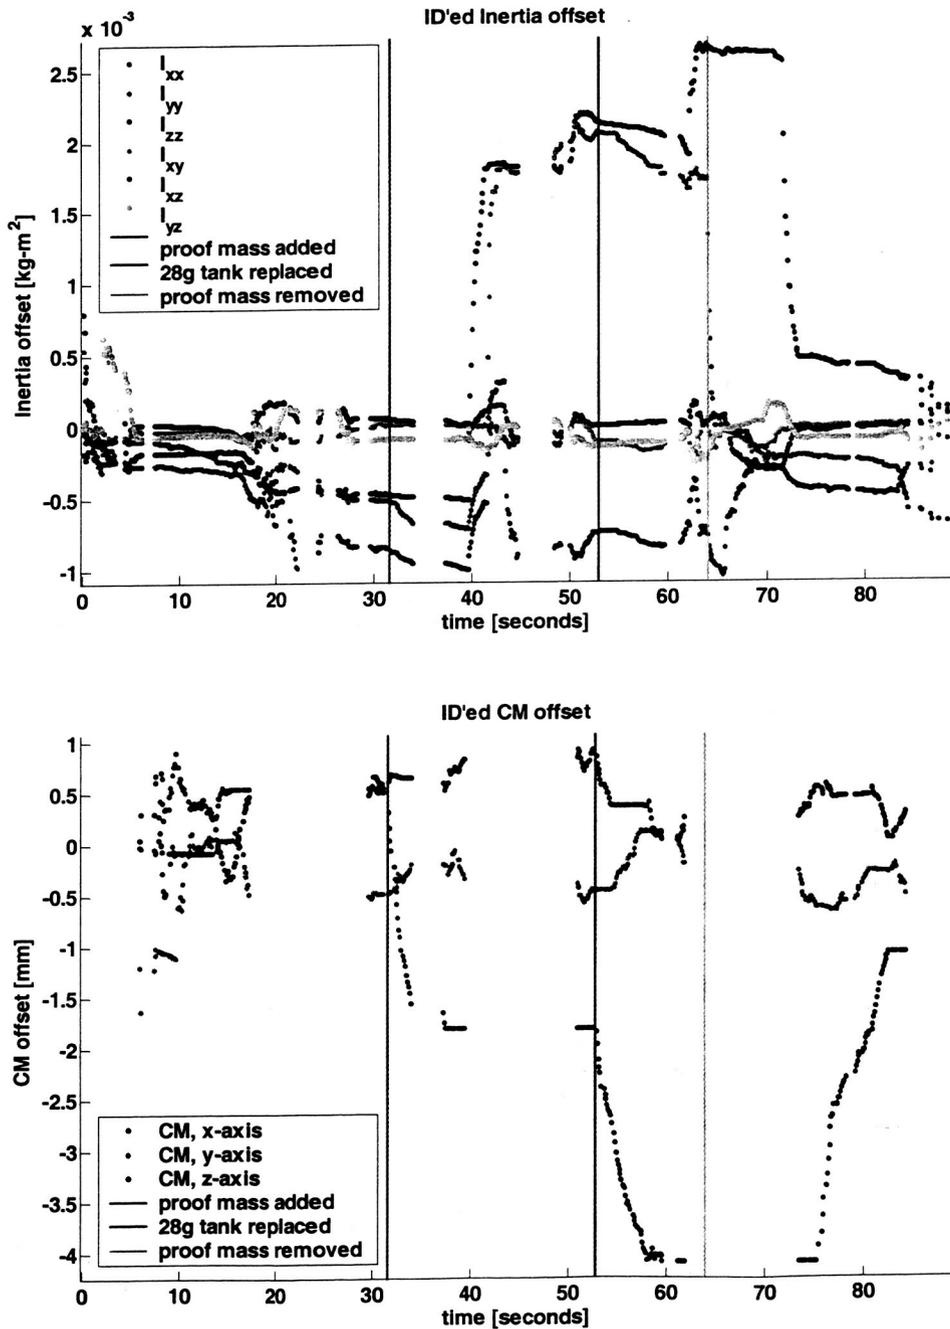
Raw gyro integers, sampled at 1 kHz, were also downloaded during both paused and un-paused periods. This figure shows an example of the ringing induced whenever a thruster is opened or closed. In this case, the end of an 800 ms firing period occurs at 63.0 seconds. Prior to this, the signal is fairly clean, with a negative slope due to the applied torque. A latency of a few milliseconds follows 63.0, at which point the thrusters close for a 200 ms coasting period.

This closing of the thrusters excites a 338 Hz resonance in the gyro⁵, as happens for thruster openings as well. This ringing is close to the 500 Hz Nyquist frequency (a single-pole 300 Hz analog pre-filter is used, which does not do much for this ringing), with the near-aliasing resulting in the pattern shown. The ringing is handled very effectively by the estimation implemented on-board, as seen by the relatively consistent acceleration estimates in the

⁵ "The 338 Hz oscillation is result of an acceleration input in that frequency. The rate sensor is sensitive to inputs at that frequency. 338 Hz is the frequency difference between the drive tines (of the quartz tuning fork) and the pickup tines. If there is acceleration input at [this frequency], you will see an output [at that frequency]." [Loggins 2003]

difference is attributed to a general "disturbance," even though it may be due to thruster variability, acceleration-estimation error, etc. Any sample containing a disturbance vector exceeding a certain length was judged to have been disturbed by handling. These data points were not used in the Mass ID updates, but the thruster firing data was used to update the burn-time integration (BTI) calculation that estimates the fuel mass.

These histograms show the detected "disturbances" (deviations from nominal) for all rotational axes. In the "All data" column, the rightmost bar indicates all data in that bin and beyond, most of which are very far beyond. These plots show a fairly clear distinction between disturbed and un-disturbed flight. To be safe, in places where a disturbance was detected, neighboring samples were also removed, since they may have been disturbed, just not enough to clear the threshold. The un-disturbed data is surprisingly clean (as compared to the noise in the X-38 disturbing acceleration calculations, for example), meaning that thruster FDI for this application should be straightforward. For example, the angular acceleration due to a single thruster is approximately 0.5-0.65 rad/sec², so the noise is sufficiently low to almost enable FDI based on a single detected deviation.



These are the results of identification on the cleaned data. No data is shown for times when the IDs do not update. For inertia ID, this is only during times of disturbances or gyro saturation. For CM ID, this is during coasting, as well as other conditions when it is judged that the SNR will be low. Initial estimate error covariance and exponential window size were set to respond quickly to the several changes that took place in this time. Tuning these ID parameters, depending on the expected prior confidence in the nominal values, and the expected changes (proof mass addition, etc.) to take place, will give different results.

Due to the open loop input sequence, significant time may pass between firings targeted at ID'ing CM (translational firings) or inertia (rotational firings), resulting in relatively flat sections that may not necessarily indicate convergence—the times of apparent step changes correspond more directly with the appropriate thruster firings rather than when the mass changes were made.

The values shown are deviations from the nominal mass properties. These "nominal" properties were in fact generated from the analysis of all Mass ID data sets with no proof mass from the November 2003 KC flights (14 parabolas), since it was judged that they probably represented the best data available. The ID aims to identify the deviations of the dry vehicle properties, assuming the BTI calculation of fuel mass properties is done exactly.

The identified proof mass-induced change in mass properties was about 80-85% of that expected based on inertia calculations. This error was observed consistently across all data sets from the Nov 2003 KC tests. The error can be partially attributed to un-resolved thruster characterization issues, and potentially inaccurate measurement of the proof mass location with respect to the geometric center. It will be investigated further prior to space testing.

These are the nominal (and identified) values for the inertia (about CM) and CM (in coordinates centered at the geometric center of the vehicle) of SPHERES #5 with an empty fuel tank:

$$I = \begin{bmatrix} 0.022039 & 0.000196 & -0.000055 \\ 0.000196 & 0.019684 & -0.000215 \\ -0.000055 & -0.000215 & 0.018165 \end{bmatrix} \text{ kg-m}^2, \quad CM = \begin{bmatrix} -0.016 \\ -0.821 \\ 3.082 \end{bmatrix} \text{ mm}$$

SUMMARY AND CONCLUSIONS

Algorithms for thruster FDI and mass-property identification that can run using gyro signals only have been presented. The maximum-likelihood-based thruster FDI algorithm is capable of reliably detecting and identifying hard, abrupt single- and multiple-jet on- or off-failures. The computationally efficient identification algorithms reliably and accurately identify mass properties (inertia, inverse-inertia, center of mass) in the presence of several significant noise sources.

Implementation on the MIT SPHERES experimental spacecraft and zero-g flight testing aboard the NASA zero-g KC-135A aircraft have demonstrated their viability for on-board implementation. The flight data presented here indicates that the on-board rate/acceleration estimator is providing accurate estimates despite thruster-induced gyro ringing. Some remaining thruster characterization is required before the numerical estimates can be considered accurate, but this data validates the performance of the mass identification algorithms and their real-time implementation. The much longer flight times expected on the ISS should enable excellent identification. Fuel slosh does not appear to be noticeable above the noise created by the gyro ringing and thruster variability; the effect of the fuel mass moving about averages out fairly well. Response to changed mass properties is shown to be fast enough to consider updating vehicle properties for a servicing spacecraft that may have docked with another spacecraft or added or removed a payload.

Since these technologies use existing sensors, the improved system robustness and performance that comes with the thruster fault tolerance and system identification can be achieved through a software-only implementation. Although widely applicable, these technologies appear to be especially well suited for small, unmanned, maneuvering spacecraft that are subjected to significant mass-property uncertainty, due either to fuel consumption, internal reconfiguration, or the carrying of variable payloads.

ACKNOWLEDGMENTS

Research funded by NASA Headquarters, HQ AA, PWC 349-00, and by the NASA Intelligent Systems Program 302-10-10 (part of the CICT Program). Thanks to Prof. David Miller, Steve Sell, Alvar Saenz-Otero, Mark Hilstad, Simon Nolet, and other members of the MIT/PSI SPHERES team for integration and experimental support. Thanks to Rodolfo Gonzalez, Dr. Steven Fredrickson, and Tim Straube of NASA Johnson Space Center and Dave Hammen of LinCom Corp. for providing valuable insights and information for the X-38 and Mini-AERCAM FDI applications. Thanks to William Readdy and Gary Martin of NASA headquarters for programmatic support.

REFERENCES

[AERCAM 2003] <http://aercam.nasa.gov>

- [Bergmann 1987] Bergmann, E.V., Walker, B.K., and Levy, D.R., Mass property estimation for control of asymmetrical satellites, *Journal of Guidance, Control, and Dynamics*, Vol. 10, Sept.-Oct. 1987, p. 483-491.
- [Bergmann 1990] Bergmann, E.V. and Dzielski, J., Spacecraft mass property identification with torque-generating control, *Journal of Guidance, Control, and Dynamics*, Vol. 13, Jan-Feb. 1990, p. 99-103.
- [Berkovitz 2003] Berkovitz, D., Kong, E., and Miller, D., System identification of the SPHERES autonomous rendezvous and docking testbed. *AIAA Space 2003*, Long Beach, California, 23-25 September, 2003.
- [Deyst 1976] Deyst, J. J. and Deckert, J. C., Maximum likelihood failure detection techniques applied to the Shuttle RCS jets. *Journal of Spacecraft and Rockets*, vol. 13, no. 2, 65-74, February 1976.
- [Gelb 1974] Gelb, A, *et al.*, *Applied Optimal Estimation*, MIT Press, Cambridge, Mass., 1974.
- [Isermann 1984] Isermann, R., Process Fault Detection Based on Modeling and Estimation Methods – A Survey, *Automatica*, Vol. 20, No. 4, pp. 387-404, 1984.
- [Lee 1998] Lee, A.Y. and Brown, M.J. A model-based thruster leakage monitor for the Cassini spacecraft. In *Proceedings of the American Control Conference*, 1998, vol. 2, pp 902-904, 1998.
- [Loggins 2003] Loggins, Brandon. E-mail communication to Edward Wilson. Systron Donner Inertial Division, BEI Technologies, Inc. Concord, CA. 24 February 2003.
- [MATLAB 2003] MATLAB is a registered trademark of The MathWorks, Inc., 24 Prime Park Way, Natick, MA 01760, 508-647-7000.
- [Richfield 1988] Richfield, R.F., Walker, B.K., and Bergmann, E.V., Input selection for a second-order mass property estimator, *Journal of Guidance, Control, and Dynamics*, Vol. 11, May-June 1988, p. 207-212.
- [SPHERES 2003] <http://ssl.mit.edu/spheres/>
- [Tanygin 1997] Tanygin, S. and Williams, T., Mass property estimation using coasting maneuvers, *Journal of Guidance, Control, and Dynamics*, Vol. 20, Jul-Aug 1997, p. 625-632.
- [Wilson 1995a]⁶ Wilson, E. and Rock, S.M., Reconfigurable control of a free-flying space robot using neural networks. In *Proceedings of the 1995 American Control Conference*, Seattle WA, June 1995.
- [Wilson 1995b] Wilson, E., *Experiments in Neural Network Control of a Free-Flying Space Robot*. PhD thesis, Stanford University, Stanford, CA 94305, March 1995.
- [Wilson 2002a] Wilson, E., Lages, C., and Mah, R., Gyro-based maximum-likelihood thruster fault detection and identification. In *Proceedings of the 2002 American Control Conference*, Anchorage, AK, May 2002.
- [Wilson 2002b] Wilson, E., Lages, C., and Mah, R., On-line, gyro-based, mass-property identification for thruster-controlled spacecraft. In *Proceedings of the 2002 IEEE International Midwest Symposium on Circuits and Systems*, Tulsa, OK, Aug 2002.
- [X38 2003] <http://spacelink.nasa.gov/NASA.Projects/Aerospace.Technology/Research.Aircraft/X-38.Crew.Return.Vehicle/index.html>

⁶ This and selected other references are currently available at <http://intellization.com/files/>